arXiv:1610.03627v1 [q-bio.NC] 12 Oct 2016

Computing Integrated Information

Stephan Krohn[1] and Dirk Ostwald[1,2]

[1]Computational Cognitive Neuroscience Laboratory
Department of Education and Psychology, Freie Universität Berlin
[2]Center for Adaptive Rationality
Max-Planck Institute for Human Development, Berlin

Abstract

Integrated information theory has received considerable attention as a novel research paradigm for the study of consciousness in recent years. Integrated information theory essentially proposes that quantitative consciousness is identical to integrated information, quantified by a measure called ϕ. However, this measure has mainly been developed in discrete state systems, impeding its computation based on experimental data. In the current work, based on the most recent instantiation of the theory ("Integrated information theory 3.0"), we make progress towards overcoming this limitation by making the following novel contributions: First, we formulate ϕ with respect to the general language of probabilistic models. Second, we derive a constructive rule for the decomposition of a system into two disjoint subsystems, central to the definition of ϕ. Third, we validate this general formulation by evaluating ϕ in a previously established discrete state example system. Fourth, we show that the general formulation improves parsimony, as it eschews undefined conditional distributions and the need for "virtual elements". Fifth, we transfer ϕ to the context of continuous linear Gaussian systems, a versatile model class suitable for the modelling of experimental data. Finally, by evaluating ϕ in parallel with the system state evolution, we explore the temporal dynamics of integrated information in the different model classes. In sum, the present formulation offers the prospect of computing integrated information from experimental data and may serve as a starting point for the further development of integrated information theory in an explicit probabilistic framework.

1 Introduction

Integrated information theory (Tononi, 2004, 2005, 2008, 2012; Oizumi et al., 2014; Tononi, 2015; Tononi et al., 2016) has recently received considerable attention as a novel research paradigm for the study of the physical substrates of consciousness. Integrated information theory essentially proposes that quantitative consciousness, i.e., the degree to which a system is conscious, is identical to its state-dependent level of integrated information, which can be quantified in a measure called "ϕ". Integration here means that the information generated by the system as a whole is in some measurable sense more than the information generated by its constituent elements and intuitively corresponds to finding

an index of the system's functional irreducibility. While this approach is not undisputed (e.g., Aaronson, 2014; Cerullo, 2015), the idea of measuring integrated information has by now gained widespread popularity in the cognitive neuroscience literature and beyond (e.g., Balduzzi and Tononi, 2008, 2009; Deco et al., 2015; Koch et al., 2016; Tegmark, 2016).

Ultimately, the aim of integrated information theory is to evaluate a measure of quantitative consciousness based on empirical data, such as electroencephalographic (EEG) or functional magnetic resonance imaging (fMRI) recordings. However, to date the theory underlying the integrated information measure ϕ is primarily developed with respect to specific examples, usually low-dimensional discrete state systems that implement logical operations. For the computation of ϕ for a wide variety of systems and based on different data types, a sufficiently general formulation of ϕ is required. While there have been some recent efforts to make the computation of ϕ feasible for empirical data sets (Barrett and Seth, 2011; Oizumi et al., 2016; Tegmark, 2016), no such effort exists for the latest instantiation of integrated information theory, developed in Oizumi et al. (2014) and referred to as "integrated information theory 3.0".

With the current work, we aim to achieve a step in the direction of the experimental applicability of integrated information theory 3.0. Our approach is based on the fundamental idea that the evaluation of integrated information from empirical data can be achieved by estimating the parameters of a probabilistic model from the data, and applying the measure ϕ to the thus estimated system model (e.g., Ostwald et al., 2010, 2014; Ostwald and Starke, 2016). To compute ϕ for such a model, we here make the following novel contributions: First, we formulate ϕ with respect to the general language of probabilistic models, by which we mean joint probability distributions over random entities (e.g., Efron and Hastie, 2016; Gelman et al., 2014; Barber, 2012; Murphy, 2012). Second, we derive a constructive rule for the decomposition of a system into two disjoint subsystems, central to the definition of ϕ. Third, we validate this general formulation by evaluating ϕ in a previously established discrete state example system. Fourth, we show that the general formulation improves parsimony, as it eschews undefined conditional probability distributions and the need for "virtual elements" as introduced by Oizumi et al. (2014). Fifth, we transfer ϕ to the context of continuous linear Gaussian systems (LGS), a versatile model class suitable for the modelling of real-life experimental data (e.g., Roweis and Ghahramani, 1999). Finally, by evaluating ϕ in parallel with the system state evolution, we explore the temporal dynamics of integrated information in the different model classes. With respect to the refinements of integrated information theory 3.0, we limit ourselves to the treatment of ϕ ("small phi") throughout. We eschew a detailed discussion of Φ ("capital phi"), because the development of the latter essentially corresponds to a reiteration of the framework for ϕ over all possible subsets of a system. For a detailed differentiation of the various "phi" measures proposed in the context of integrated information theory 3.0 and their relation to the ϕ discussed herein, please see Appendix A. Note that henceforth, we shall use the abbreviation "IIT" to refer exclusively to integrated information theory 3.0 as developed in Oizumi et al. (2014), unless noted otherwise.

We proceed as follows. In Section 2 *Computing ϕ : Theory*, we first discuss the scope of integrated information and provide an explicit technical definition of ϕ in the subsections *System model* and *Definition of ϕ*. In a hierarchical fash-

ion, we then detail the conceptual and technical underpinnings of the constituent elements of this definition in the subsections *Effect and cause repertoires* and *System decomposition*. Having established a general definition of ϕ, we then put this definition to work in the Section 3 *Computing ϕ : Applications*. To validate our approach, we first apply the general framework in a discrete state system setting, rooted in the development of IIT (Oizumi et al., 2014). Subsequently, we apply the same general framework to the model class of linear Gaussian systems and implement the evaluation of ϕ in Gaussian example systems. Finally, we discuss our approach for computing ϕ with respect to related work in the literature and potential future developments.

Notation and Implementation

A few remarks on our notation of probabilistic concepts are in order. To denote random variables/vectors and their probability distributions, we use an applied notation throughout. This means that we eschew a discussion of a random entity's underlying measure-theoretic probability space model (e.g., Billingsley, 2008), and focus on the random entity's outcome space and probability distribution. For a random variable/vector X, we denote its distribution by $p(X)$, implicitly assuming that this may be represented either by a probability mass or a probability density function. To denote different distributions of the same random variable/vector, we employ subscripts. For example, $p_a(X)$ is to indicate a probability distribution of X that is different from another probability distribution $p_b(X)$. In the development of ϕ, stochastic (conditional) dependencies between random variables are central. To this end, we use the common notation that the statement $p(X|Y) = p(X)$ is meant to indicate the stochastic independence of X from Y and the statement $p(X|Y, Z) = p(X|Z)$ is meant to indicate the (stochastic) conditional independence of X on Y given Z (Dawid, 1979; Geiger et al., 1990). Finally, we use $N(X; \mu, \Sigma)$ to denote the Gaussian probability density function of a d-dimensional random vector X with expectation parameter $\mu \in \mathbb{R}^d$ and positive definite covariance matrix parameter $\Sigma \in \mathbb{R}^{d \times d}$, I_d denotes the $d \times d$ identity matrix, and 1_d denotes a d-dimensional vector of all ones. All Matlab code (The MathWorks, Inc., Natick, MA, United States) developed to implement our formulation of IIT and generate the technical figures herein is available from the Open Science Framework (`https://osf.io/hb4a5/`).

2 Computing ϕ : Theory

System model

IIT models the temporal evolution of a system by a discrete time multivariate stochastic process (Cox and Miller, 1977)

$$p(X_1, X_2, ..., X_T). \tag{1}$$

In the probabilistic model (1), $X_t, t = 1, ..., T$ denotes a finite set of d-dimensional random vectors. Here, the limitation to a finite set of discrete time-points is primarily motivated by the eventual goal to apply the concepts of IIT in a data-analytical setting, not by inherent constraints of IIT. Each random vector X_t

comprises random variables $x_{t_i}, i = 1, 2, ..., d$ ($d \in \mathbb{N}$) that may take on values in one-dimensional outcome spaces $\mathcal{X}_1, \mathcal{X}_2 ..., \mathcal{X}_d$, such that

$$X_t = (x_{t_1}, x_{t_2}, ..., x_{t_d})^T \tag{2}$$

may take on values in the d-dimensional outcome space $\mathcal{X} := \prod_{i=1}^{d} \mathcal{X}_i$. We assume $\mathcal{X} \subseteq \mathbb{R}^d$ throughout. IIT further assumes that the stochastic process fulfils the Markov property, i.e., that the probabilistic model (1) factorizes according to

$$p(X_1, X_2, ..., X_T) = p(X_1) \prod_{t=2}^{T} p(X_t | X_{t-1}), \tag{3}$$

and that the ensuing Markov process is time-invariant, i.e. that all conditional probability distributions $p(X_t|X_{t-1})$ on the right-hand side of eq. (3) are identical (Figure 1A). We will refer to $p(X_t|X_{t-1})$ as the system's *transition probability distribution* in the following. Finally, IIT assumes that the random variables constituting X_t are conditionally independent given X_{t-1}, i.e., that the conditional distribution $p(X_t|X_{t-1})$ factorizes according to

$$p(x_{t_1}, x_{t_2}, ..., x_{t_d} | X_{t-1}) = \prod_{i=1}^{d} p(x_{t_i} | X_{t-1}). \tag{4}$$

Definition of ϕ

Based on the assumptions of eqs. (1), (3) and (4), IIT defines the integrated information measure ϕ of a system state $X \in \mathcal{X}$ as follows:

$$\phi : \mathcal{X} \to \mathbb{R}, X \mapsto \phi(X) := \min \left\{ \phi_e(X), \phi_c(X) \right\}, \tag{5}$$

where $\phi_e : \mathcal{X} \to \mathbb{R}$ and $\phi_c : \mathcal{X} \to \mathbb{R}$ are defined as

$$\phi_e(X) := \min_{i \in I} \left\{ D \left(p_e(X_t | X_{t-1} = X) || p_e^{(i)}(X_t | X_{t-1} = X) \right) \right\} \tag{6}$$

and

$$\phi_c(X) := \min_{i \in I} \left\{ D \left(p_c(X_{t-1} | X_t = X) || p_c^{(i)}(X_{t-1} | X_t = X) \right) \right\}, \tag{7}$$

respectively. In (7) and (6),

- $\phi_e(X)$ and $\phi_c(X)$ are referred to as *integrated effect information* and *integrated cause information* of the system state X,
- $p_c(X_{t-1}|X_t = X)$ and $p_e(X_t|X_{t-1} = X)$ are conditional probability distributions that are constructed from the transition probability $p(X_t|X_{t-1})$ of the stochastic process as detailed below and are referred to as the *effect repertoire* and *cause repertoire* of X, respectively,
- $p_e^{(i)}(X_t|X_{t-1} = X)$ and $p_c^{(i)}(X_{t-1}|X_t = X)$ are "decomposed variants" of the effect and cause repertoires, that result from the removal of potential stochastic dependencies in the system's transition probability distribution as detailed below,

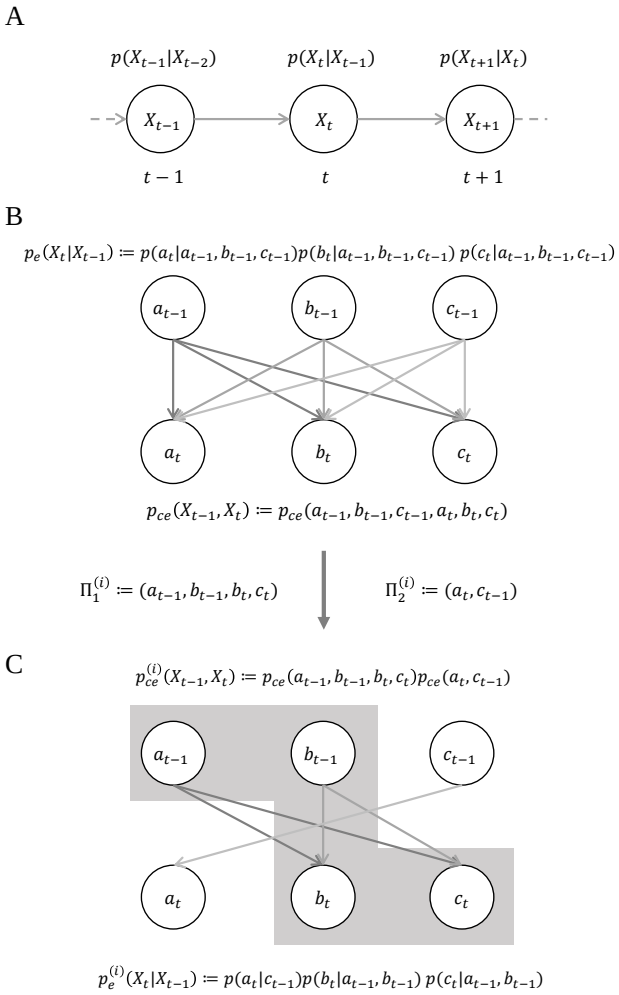

Figure 1: System model and system decomposition. IIT models the system of interest by a time-invariant first-order Markov process, depicted in Panel A as graphical model (e.g., Bishop, 2006). Nodes denote random vectors and directed links denote the stochastic dependence of the child node on the parent node. Panels B and C display the exemplary decomposition of a three-dimensional system with state vector $X_t := (a_t, b_t, c_t)$ as graphical model. Here, nodes denote the constituent random variables variables of the random vectors X_{t-1} and X_t. Panel B depicts the unpartitioned system, in which all potential stochastic dependencies of the elements are visualized. The constituent random variables of X_t are conditionally independent given X_{t-1} (cf. eq. (4)), and the joint distribution $p_{ce}(X_{t-1}, X_t)$ is invoked by the assumption of an uncertain marginal distribution $p_u(X_t)$ for each $t = 2, ..., T$. Panel C shows an exemplary decomposition of the system, which is based on the bipartition of (X_{t-1}, X_t) into the subsets $\Pi_1^{(i)} = \{a_{t-1}, b_{t-1}, b_t, c_t\}$ (gray inset) and $\Pi_2^{(i)} = \{a_t, c_{t-1}\}$. In the factorized joint distribution $p_{ce}(X_{t-1}, X_t)$ the directed links across the partition boundary are removed, while the links within each partition remain.

- I is an index set, the elements of which index the "decomposed variants" of the effect and cause repertoires, and
- $D : P \times P \to \mathbb{R}_+, (p_1, p_2) \mapsto D(p_1||p_2)$ denotes a divergence measure between (conditional) probability distributions over the same random entity, with P indicating the set of all possible distributions of this entity. In practice, D corresponds to the the earth mover's distance for discrete state systems (Levina and Bickel, 2001; Mallows, 1972), and the Kullback-Leibler divergence (Kullback and Leibler, 1951) for continuous state systems.

We next discuss the intuitive and technical underpinnings of the constituents of the definition ϕ by eqs. (5) to (7) in further detail.

Effect and cause repertoires

Intuitively, IIT aims to quantify how a given system state constrains both future and previous system states. Formally, this is achieved by constructing the *effect* and *cause repertoire* conditional probability distributions as follows.

The *effect repertoire* is defined as the conditional distribution of X_t given X_{t-1} and thus corresponds to the stochastic process' transition probability distribution

$$p_e(X_t|X_{t-1}) := p(X_t|X_{t-1}). \tag{8}$$

To construct the *cause repertoire*, IIT implicitly defines a joint distribution p_{ce} over X_{t-1} and X_t by multiplication of the Markov transition probability distribution $p(X_t|X_{t-1})$ with a marginal distribution $p_u(X_{t-1})$, i.e.

$$p_{ce}(X_{t-1}, X_t) := p_u(X_{t-1})p(X_t|X_{t-1}). \tag{9}$$

Here, the marginal distribution $p_u(X_{t-1})$ is meant to represent a maximum of uncertainty about X_{t-1}, and for the case of a finite outcome space $\mathcal{X}$, corresponds to the uniform distribution over all states. Based on the joint distribution of (9), the cause repertoire is then defined as the conditional distribution of X_{t-1} given X_t:

$$p_c(X_{t-1}|X_t) := \frac{p_{ce}(X_{t-1}, X_t)}{\sum_{X_{t-1}} p_{ce}(X_{t-1}, X_t)} \tag{10}$$

Note that also the effect repertoire may be recovered from (10) by conditioning on X_{t-1}.

System decomposition

IIT's central concept of *integrated information* is based on the intuition that integrated information is "information generated by the system above and beyond the information generated by its parts" (Oizumi et al., 2014). To evaluate integrated information, IIT first considers all possible ways to decompose a system into two subsystems that do not influence each other. IIT then aims to identify the system decomposition which, for a given system state, is most similar to the actual system in terms of the divergence between the system state's effect and cause repertoires (cf. eqs. (6) and (7)). In the following, we first specify what this intuitive notion of integration means in technical terms with respect to the

system model of eqs. (1), (3) and (4). We then present a formal rule for the evaluation of the necessary decompositions.

In technical terms, the "system" to be decomposed corresponds to the collection of random variables and their conditional dependencies that define the discrete time multivariate stochastic process (cf. eq. (1)). Because of the process' time-invariant Markov property (cf. eq. (3)), the relevant random variables are the constituents of two time-adjacent random vectors X_{t-1} and X_t. As seen above, based on an uncertain marginal distribution over X_{t-1}, one may define a joint distribution $p_{ce}(X_{t-1}, X_t)$ of these vectors for each $t = 2, ..., T$. Note that the joint distribution $p_{ce}(X_{t-1}, X_t)$ can equivalently be regarded as a joint distribution over the set of all constituent random variables of the random vectors X_{t-1} and X_t,

$$(X_{t-1}, X_t) := \{x_{t-1_1}, x_{t-1_2}, ..., x_{t-1_d}, x_{t_1}, x_{1_2}, ..., x_{t_d}\}. \tag{11}$$

IIT then uses the intuitive appeal of graphical models (Lauritzen, 1996; Jordan, 1998) to introduce the idea of "cutting a system" into two independent parts. Technically, cutting the graphical model of $p_{ce}(X_{t-1}, X_t)$ corresponds to (a) partitioning the set of random variables in eq. (11) into two disjoint subsets and (b) removing all stochastic dependencies across the boundary between the resulting random variable subsets while retaining all conditional dependencies within each subset (Figure 1B and 1C). Notably, there are $k := 2^{2d-1} - 1$ unique ways to bipartition a set of cardinality $2d$ (Appendix B). This corresponds to k ways of cutting the corresponding graphical model and thus induces a set of k differently factorized joint distributions $p^{(i)}_{ce}(X_{t-1}, X_t), i = 1, ..., k$, which form the basis for the decomposed effect and cause repertoires $p^{(i)}_e(X_t|X_{t-1})$ and $p^{(i)}_c(X_{t-1}|X_t)$ in the definition of ϕ (cf. eqs. (6) and (7)).

We next formalize the construction of $p^{(i)}_{ce}(X_{t-1}, X_t)$ for $i = 1, ..., k$. To this end, first recall that a partition of a set S is a family of sets P with the properties

$$\emptyset \notin P, \bigcup_{M \in P} M = S, \text{ and if } M, M' \in P \text{ and } M \neq M', \text{ then } M \cap M' = \emptyset. \tag{12}$$

Let $\Pi^{(i)}$ denote a bipartition of the set of random variables (X_{t-1}, X_t), i.e.

$$\Pi^{(i)} := \left(\Pi^{(i)}_1, \Pi^{(i)}_2\right), \tag{13}$$

where

$$\Pi^{(i)}_1, \Pi^{(i)}_2 \subset (X_{t-1}, X_t), \Pi^{(i)}_1 \cap \Pi^{(i)}_2 = \emptyset \text{ and } \Pi^{(i)}_1 \cup \Pi^{(i)}_2 = (X_{t-1}, X_t). \tag{14}$$

Let further

$$p_{ce}(\Pi^{(i)}_1) = \sum_{\Pi^{(i)}_2} p_{ce}(X_{t-1}, X_t) \text{ and } p_{ce}(\Pi^{(i)}_2) = \sum_{\Pi^{(i)}_1} p_{ce}(X_{t-1}, X_t) \tag{15}$$

denote the marginal distributions of $p_{ce}(X_{t-1}, X_t)$ (cf eq. (9)) of the random variables contained in $\Pi^{(i)}_1$ and $\Pi^{(i)}_1$, respectively. Then the elements of the set of factorized variants of the joint distribution $p_{ce}(X_{t-1}, X_t)$ are given by

$$p^{(i)}_{ce}(X_{t-1}, X_t) := p_{ce}(\Pi^{(i)}_1) p_{ce}(\Pi^{(i)}_2) \text{ for } i = 1, 2, .., k. \tag{16}$$

In summary, the measure of integrated information ϕ rests on a standard probabilistic model approach to dynamical systems - a multivariate stochastic process that fulfils the Markov property (cf. eqs. (1), (3) and (4)). On this background, the integrated information of a system state is defined by the minimum of the state's integrated effect and cause information. These quantities are defined as the minima of the divergences between the state conditional probability distributions (evaluated directly in terms of the system's transition probability distribution (cf. eq. (8)) and from an approximation of the joint distribution over adjacent time points (cf. eq. (9)) and variants of these distributions that result from the removal of stochastic dependencies in the system's transition probability distribution (cf. eq. (16)). In the following sections, we show how this general definition of ϕ can be applied in the context of both discrete and continuous state systems.

3 Computing ϕ : Applications

In the current section, we consider two different types of probabilistic models and evaluate ϕ in exemplary instances of each type. We first apply the theoretical development of ϕ to a system with discrete state space which is defined non-parametrically by the explicit definition of the transition probability distribution factors. This system corresponds to the exemplary system discussed in (Oizumi et al., 2014) and serves as a validation of our formulation for computing ϕ. Subsequently, we consider the class of linear Gaussian systems, which constitute a highly versatile model class in the analysis of empirical data, such as EEG or fMRI measurements.

3.1 Discrete state systems

In discrete state systems, the random variables that model the system's elements take on a finite number of states with a certain probability mass. Discrete state system constitute the main model class that has been employed throughout the development of integrated information theory (Balduzzi and Tononi, 2008, 2009; Oizumi et al., 2014). As an exemplary discrete state system, we consider a system presented in Oizumi et al. (2014). This system is three-dimensional, and, in concordance with Oizumi et al. (2014), we denote its state vector by $X_t = (a_t, b_t, c_t)$ (Figure 2A). The system is defined in terms of the marginal conditional distributions of its component variables (cf. eq. (4)). Specifically, the variables a_t, b_t and c_t may take on values in $\{0, 1\}$, such that the outcome space $\mathcal{X}$ is defined as $\{0, 1\}^3$, and implement logical operations on the state of their predecessors a_{t-1}, b_{t-1} and c_{t-1}. As shown in Figure 2B, a_t implements a logical OR, b_t implements a logical AND, and c_t implements a logical XOR operation. Note that in this case, the relevant distributions of eq. (3) correspond to probability mass functions, which can be represented on the implementational level by high-dimensional numerical arrays.

As discussed above, the effect repertoire $p_e(X_t|X_{t-1})$ of the system corresponds to the product of the marginal conditional distributions and is identical with the system's forward transition probability matrix (cf. eq. (4)). This distribution is shown in Figure 2C. The joint distribution $p_{ce}(X_{t-1}, X_t)$ is derived by multiplication of the transition probability with a maximally uncertain dis-

tribution over past states $p_u(X_{t-1})$ (cf.eq. (9)). In this example, the maximally uncertain distribution is given by the uniform distribution over past states, i.e. $p_u(X_t = X_t^*) := |\{0,1\}^3|^{-1}$ for all $X_t^* \in \{0,1\}^3$ (cf. Figure 2D). From the ensuing joint distribution $p_{ce}(X_{t-1}, X_t) = p(a_{t-1}, b_{t-1}, c_{t-1}, a_t, b_t, c_t)$, the cause repertoire $p_c(X_{t-1}|X_t)$ can be evaluated by conditioning on X_t. The resulting distribution is shown in Figure 2E. Note that there are some undefined entries (displayed in red). These undefined entries correspond to system states that cannot have been caused by any previous state due to the constraints placed by the logical operations of the system variables.

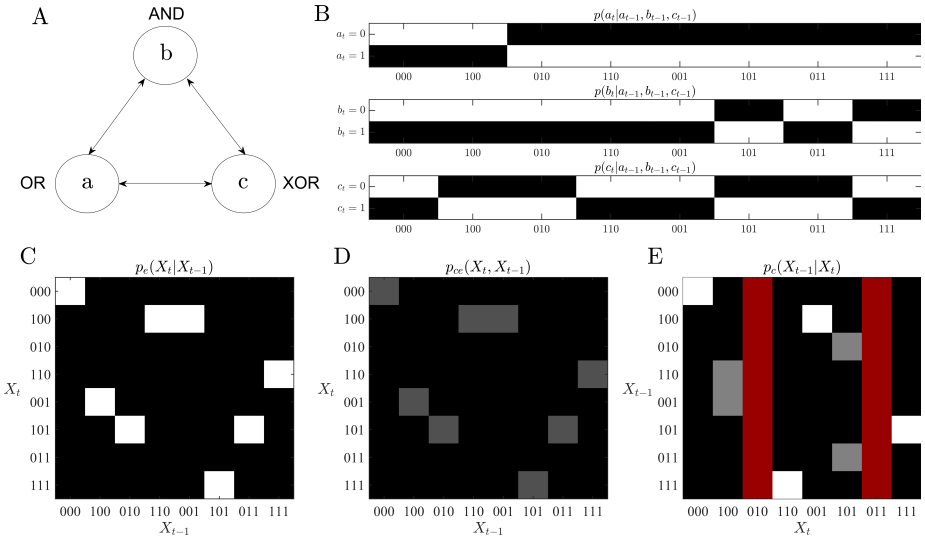

Figure 2: Evaluation of ϕ in an exemplary discrete state system. The system is identical to that presented in Oizumi et al. (2014) (e.g., Figures 1 and 4 therein). Panel A shows the system comprising three random variables that implement the logical operations OR, AND, and XOR. Panel B visualizes the corresponding marginal conditional probability distributions, with black tiles indicating a probability mass of 0 and white tiles indicating a probability mass of 1. The product of these marginal conditional probability distributions yields the conditional distribution $p(X_t|X_{t-1})$ depicted in panel C, i.e. the state transition probability matrix, which is identical to the effect repertoire $p_e(X_t|X_{t-1})$. By multiplication with a maximally uncertain distribution over past states, i.e. $p_u(X_{t-1})$, the joint distribution $p_{ce}(X_t, X_{t-1})$ of panel D is obtained. Here, dark gray tiles indicate a probability mass of 0.125. For the current example, $p_u(X_{t-1})$ corresponds to the uniform distribution over system states. The joint distribution forms the basis for all k factorizations in the calculation of ϕ. Finally, conditioning $p_{ce}(X_t, X_{t-1})$ on X_t yields the cause repertoire $p_c(X_{t-1}|X_t)$ as shown in panel E. Here, white tiles indicate a probability mass of 1, gray tiles a probability mass of 0.5, and red tiles represent undefined entries. These entries correspond to states of X_t that cannot have been caused by any of the states of X_{t-1} due to the logical structure of the network.

Decomposed effect and cause repertoires

To compute ϕ, the decomposed effect and cause repertoires associated with the factorized joint distributions $p_{ce}^{(i)}(X_{t-1}, X_t), i = 1, ..., k$ (cf. eq. (16)) have to

be evaluated by conditioning these distributions on X_{t-1} or X_t, respectively. Notably, for each bipartition of $(X_{t,1}, X_t)$ the factorization of the joint distribution $p_{ce}(X_{t-1}, X_t)$ (cf. eq. (9)) also induces a factorization of the respective marginal distributions of the constituent random variables of X_{t-1} or X_t, which needs to be accounted for. In the following, we first specify how the resulting decomposed effect and cause repertoires can be evaluated in a well-defined and general manner and then illustrate this process in the example system under consideration.

Let $\Theta \subset (X_{t-1}, X_t)$ denote the set of random variables that the factorized joint distribution $p_{ce}^{(i)}(X_{t-1}, X_t)$ (cf. eq. (16)) is to be conditioned on, i.e. $\Theta := (X_t)$ or $\Theta := (X_{t-1})$. Let further

$$p_{ce}^{(i)}(\Theta) = \sum_{(X_t - 1, X_t) \setminus \Theta} p_{ce}(X_{t-1}, X_t) \tag{17}$$

be the marginal distribution of eq. (9) of the random variables in Θ. Define

$$\Theta_1^{(i)} = \Theta \cap \Pi_1^{(i)} \text{ and } \Theta_2^{(i)} = \Theta \cap \Pi_2^{(i)}. \tag{18}$$

Note that because $\left(\Pi_1^{(i)}, \Pi_2^{(i)}\right)$ is a partition of (X_{t-1}, X_t), it follows that

$$\Theta_1^{(i)} \cup \Theta_2^{(i)} = \Theta \text{ and } \Theta_1^{(i)} \cap \Theta_2^{(i)} = \emptyset, \tag{19}$$

and that, possibly, $\Theta_1^{(i)} = \emptyset$ or $\Theta_2^{(i)} = \emptyset$. By definition, the marginal distributions of $p_{ce}(\Theta)$ with respect to $\Theta_1^{(i)}$ and $\Theta_2^{(i)}$ are given by

$$p_{ce}\left(\Theta_1^{(i)}\right) = \sum_{\Theta_2^{(i)}} p_{ce}(\Theta) \text{ and } p_{ce}\left(\Theta_2^{(i)}\right) = \sum_{\Theta_1^{(i)}} p_{ce}(\Theta). \tag{20}$$

Finally, define

$$p_{ce}(\Theta^{(i)}) := \begin{cases} p(\Theta_1^{(i)}) & \text{, if } \Theta_2^{(i)} = \emptyset \\ p(\Theta_2^{(i)}) & \text{, if } \Theta_1^{(i)} = \emptyset \\ p(\Theta_1^{(i)}) p(\Theta_2^{(i)}) & \text{, else .} \end{cases} \tag{21}$$

Based on eq. (15) and eq. (21), the decomposed effect and cause repertoires in the definition of ϕ (cf. eqs. (6) and (7)) are then defined as

$$p_e^{(i)}(X_t | X_{t-1}) = \frac{p_{ce}(\Pi_1^{(i)}) p_{ce}(\Pi_2^{(i)})}{p_{ce}(\Theta^{(i)})} \text{ for } \Theta = (X_{t-1}) \tag{22}$$

and

$$p_c^{(i)}(X_{t-1} | X_t) = \frac{p_{ce}(\Pi_1^{(i)}) p_{ce}(\Pi_2^{(i)})}{p_{ce}(\Theta^{(i)})} \text{ for } \Theta = (X_t). \tag{23}$$

Note that for the evaluation of ϕ_c and ϕ_e, the random variables constituting X_{t-1} and X_t on the right-hand side of eqs. (22) and (23) have to be set to the values defined by the system state X.

Example

To illustrate the above, we consider the exemplary system of Figure 2. Here, the concatenated state vector over two adjacent time-points is given by (cf. eq. (11))

$$(X_{t-1}, X_t) = \{a_{t-1}, b_{t-1}, c_{t-1}, a_t, b_t, c_t\} \tag{24}$$

One of the $k = 2^{6-1} - 1 = 31$ bipartitions of eq. (24) (which we label here as $i := 1$) is given by

$$\Pi_1^{(1)} = \{a_{t-1}, b_{t-1}, b_t, c_t\} \text{ and } \Pi_2^{(1)} = \{c_{t-1}, a_t\}. \tag{25}$$

Note that this corresponds to the partition depicted in Figure 1. Hence, with eq. (15)

$$p_{ce}(\Pi_1^{(1)}) = p_{ce}(a_{t-1}, b_{t-1}, b_t, c_t) \text{ and } p_{ce}(\Pi_2^{(1)}) = p_{ce}(c_{t-1}, a_t). \tag{26}$$

For the decomposed effect repertoire corresponding to the bipartition as given in eq. (25), we have $\Theta = X_{t-1}$, and thus obtain

$$\Theta_1^{(1)} = \{a_{t-1}, b_{t-1}\} \text{ and } \Theta_2^{(1)} = \{c_{t-1}\} \tag{27}$$

Because $\Theta_1^{(1)} \neq \emptyset$ and $\Theta_2^{(1)} \neq \emptyset$, we have (cf. eq. (20) and eq. (21))

$$p_{ce}(\Theta^{(1)}) = p_{ce}(a_{t-1}, b_{t-1}) p_{ce}(c_{t-1}), \tag{28}$$

and finally

$$\begin{aligned} p_e^{(1)}(X_t | X_{t-1}) &= \frac{p_{ce}(a_{t-1}, b_{t-1}, b_t, c_t) p_{ce}(c_{t-1}, a_t)}{p_{ce}(a_{t-1}, b_{t-1}) p_{ce}(c_{t-1})} \\ &= p_{ce}(b_t | a_{t-1}, b_{t-1}) p_{ce}(c_t | a_{t-1}, b_{t-1}) p_{ce}(a_t | c_{t-1}) \end{aligned} . \tag{29}$$

Further, for the decomposed cause repertoire resulting from eq. (25), we have $\Theta = (X_t)$, and thus obtain with eq. (18)

$$\Theta_1^{(1)} = \{b_t, c_t\} \text{ and } \Theta_2^{(1)} = \{a_t\}. \tag{30}$$

Because $\Theta_1^{(1)} \neq \emptyset$ and $\Theta_2^{(1)} \neq \emptyset$, we have

$$p_{ce}(\Theta^{(1)}) = p_{ce}(b_t, c_t) p_{ce}(a_t) \tag{31}$$

and finally

$$\begin{aligned} p_c^{(1)}(X_{t-1} | X_t) &= \frac{p_{ce}(a_{t-1}, b_{t-1}, b_t, c_t) p_{ce}(c_{t-1}, a_t)}{p_{ce}(b_t, c_t) p_{ce}(a_t)} \\ &= p_{ce}(a_{t-1}, b_{t-1} | b_t, c_t) p_{ce}(c_{t-1} | a_t). \end{aligned} \tag{32}$$

For further illustration of this constructive process, a more exhaustive example is provided in the Supporting Information (Figures S1 and S2). There, the effect and cause repertoires corresponding to the seven bipartitions of a two-dimensional system are considered in detail.

Computation of ϕ

Based on the above, we are now in the position to compute ϕ for arbitrary system states in the exemplary system. As an example, Figure 3 depicts the the evaluation of $\phi(X)$ for the system state $X = (1, 0, 0)$. Specifically, Figure 3A depicts the probability mass functions underlying the evaluation of integrated effect information $\phi_e((1, 0, 0))$ and Figure 3B depicts the probability mass functions underlying the evaluation the integrated cause information $\phi_c((1, 0, 0))$. The right panel of 3A depicts the decomposed variant of the effect repertoire which results in the minimum EMD $\phi_e((1, 0, 0)) = 0.25$ between the original and decomposed effect repertoire variants (cf. eq. (6)). Notably, for the current system there are $2^{6-1} - 1 = 31$ possible ways to bipartition the graphical model of $p_{ce}(X_{t-1}, X_t)$, and hence 31 versions of $p_e^{(i)}(X_t|X_{t-1})$. For the system state $X = (1, 0, 0)$, the ensuing partitions and distributions are shown in Figure S3. Likewise, the left panel of Figure 3B depicts the cause repertoire of the system state $X = (1, 0, 0)$, which directly follows from the distribution shown in Figure 2E. The right panel of 3B depicts a decomposed variant of the cause repertoire which results in a minimum EMD of $\phi_c((1, 0, 0)) = 0.5$ between the original and decomposed effect repertoire variants (cf. eq. (7)). As shown in Figure S4, there are in fact seven partitions and ensuing distributions, for which this minimum EMD is attained. Evaluating the minimum (cf. eq. (5)) then results in an integrated information of

$$\phi((1, 0, 0)) = 0.25 \tag{33}$$

for the current system and system state of interest. Notably, this value is identical to that reported in Figure 6 of Oizumi et al. (2014).

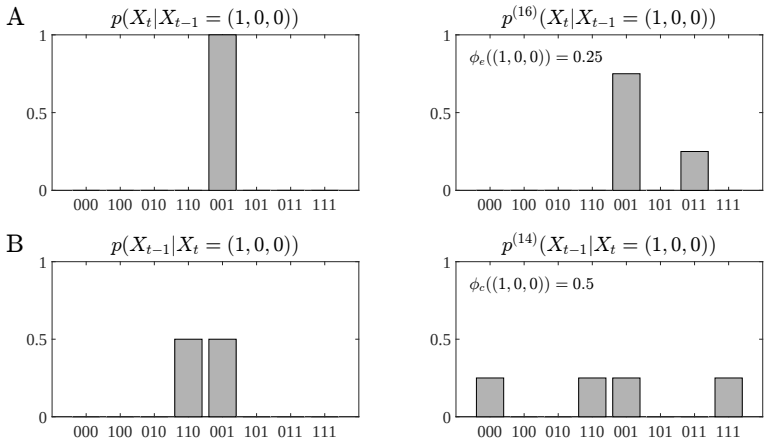

Figure 3: State-specific evaluation of ϕ in an exemplary discrete state system. The Figure depicts the probability mass functions required for the evaluation of ϕ for the system state $X = (1, 0, 0)$ in the discrete state system introduced in Figure 2. Panel A viualizes the relevant distributions required for the evaluation of $\phi_e((1, 0, 0))$ (cf. eq. (6)) and Panel B visualizes the relevant distributions required for the evaluation of $\phi_c((1, 0, 0))$ (cf. eq. (7)). The minimum of these two values is defined as the integrated information $\phi((1, 0, 0)) = 0.25$ (cf. eq. (5)) of the system state.

As a more exhaustive example, we considered the behaviour of ϕ as the exemplary system evolves over time. To this end, we precomputed ϕ for each of the eight system states of the example system, and monitored the state evolution for each possible initial state $X_1 \in \mathcal{X}$ for $t = 1, ..., 30$. Figure 4 summarizes the results of this simulation. Panel A depicts the evolution of the system state as a function of the initial state (subpanels). Notably, after less than 10 time steps, the system is either locked in an all-off state (for initial state $X_1 = (0, 0, 0)$) or oscillates between states $(1, 0, 0)$ and $(0, 0, 1)$ (for all other initial states). Panel B depicts the corresponding evaluation of $\phi(X_t)$. Mirroring the temporal evolution of the system state, $\phi(X_t)$ takes on the value 0 for all time points for initial state $X_1 = (0, 0, 0)$ or oscillates between 0.25 and 0 for all other initial states.

On empty conditional distributions and virtual elements

In contrast to the approach introduced of Oizumi et al. (2014), we consider the decomposition of a system at the level of the joint distribution $p_{ce}(X_{t-1}, X_t)$ and not at the level of the effect and cause repertoires $p_e(X_t|X_{t-1})$ and $p_c(X_{t-1}|X_t)$. In this section, we discuss the advantages of our formulation for computing ϕ and relate it to the approach discussed in Oizumi et al. (2014). To this end, we consider a fixed bipartition $\Pi = (\Pi_1, \Pi_2)$ of (X_{t-1}, X_t) (cf. eq. (13)), where we omit the superscript (i) for notational convenience. We further use

$$\Pi^1_{t-1} := X_{t-1} \cap \Pi_1 \text{ and } \Pi^2_{t-1} := X_{t-1} \cap \Pi_2 \tag{34}$$

to denote the elements of X_{t-1} that are part of the bipartition subsets Π_1 and Π_2, respectively, and

$$\Pi^1_t := X_t \cap \Pi_1 \text{ and } \Pi^2_t := X_t \cap \Pi_2 \tag{35}$$

to denote the elements of X_t that are part of the bipartition subsets Π_1 and Π_2, respectively.

Firstly, Oizumi et al. (2014) consider the decompositions of the effect and cause repertoire in the forms

$$p^{(i)}_e(X_t|X_{t-1}) := p(\Pi^1_t|\Pi^1_{t-1})p(\Pi^2_t|\Pi^2_{t-1}) \tag{36}$$

and

$$p^{(i)}_c(X_{t-1}|X_t) := p(\Pi^1_{t-1}|\Pi^1_t)p(\Pi^2_{t-1}|\Pi^2_t) \tag{37}$$

(see for example eq. (13) in the Supplementary Methods of Oizumi et al. (2014)). However, because depending on the bipartition any of the sets Π^1_t, Π^2_t, Π^1_{t-1}, or Π^2_{t-1} may be empty (consider for example the case of of a two-dimensional random vector X_t and a bipartition of the form $\Pi_1 = \{x_{t-1_1}\}$ and $\Pi_2 = \{x_{t-1_2}, x_{t_1}, x_{t_2}\}$), this approach can result in undefined system decompositions. Our formulation eschews this, because here only non-empty distributions are multiplied (cf. eq. (12)) and conditioned on (cf. eq. (21)).

Secondly, Oizumi et al. (2014) use the (non-probabilistic) concept of "virtual elements" to induce a maximal uncertain distribution over the system state X_{t-1}. To this end, the authors define random variable sets

$$X^{V1}_{t-1} := X_{t-1} \setminus \Pi^1_{t-1} \text{ and } X^{V2}_{t-1} := X_{t-1} \setminus \Pi^2_{t-1} \tag{38}$$

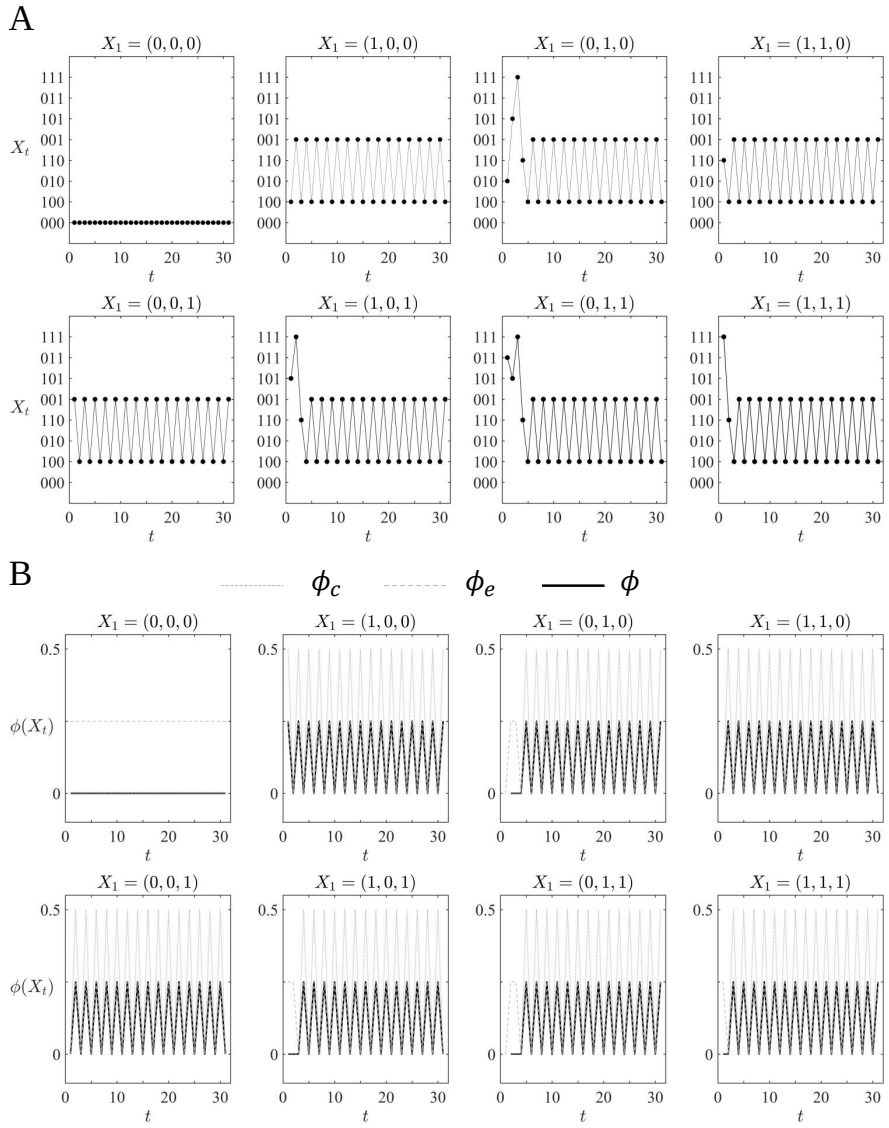

Figure 4: Evaluation of the temporal dynamics of ϕ in the discrete state example system.Panel A shows the evolution of the example system over $T = 30$ time points based on all possible states of X_1. After less than 10 time steps, the system is either locked in an all-off state (for $X_1 = (0,0,0)$) or oscillates between system states $(1,0,0)$ and $(0,0,1)$ (for all other states). Panel B shows the parallel evaluation of $\phi(X_t)$, again for all possible states of X_1.

and assume a uniform distribution $p_u(X^{V1}_{t-1}, X^{V2}_{t-1})$ when evaluating the cause and effect repertoires and their decompositions (see for example Figure 1 in the Supplementary Methods of Oizumi et al. (2014)). However, because

$$(X^{V1}_{t-1}, X^{V2}_{t-1}) = (\Pi^2_{t-1}, \Pi^1_{t-1}) = X_{t-1} \tag{39}$$

this assumption is equivalent to the assumption of eq. (9) in our formulation, which follows standard probabilistic reasoning.

3.2 Linear Gaussian Systems

As a further application of the general framework for the evaluation of ϕ, we consider the broad class of linear Gaussian systems (LGSs) (Roweis and Ghahramani, 1999). In contrast to the discrete state systems discussed above, LGSs are directly applicable to the continuous data recordings of brain imaging techniques such as EEG or fMRI. Evaluating ϕ in these kinds of probabilistic models thus paves the way for the computation of ϕ based on experimental data sets.

LGSs are basic building blocks of a wide variety of standard models in signal processing and machine learning and take the general form (cf. eq. (1) and eq. (3))

$$p(X_1, X_2, ..., X_T) = p(X_1) \prod_{t=2}^{T} p(X_t|X_{t-1}), \tag{40}$$

where

$$p(X_t|X_{t-1}) = N(X_t; AX_{t-1}, \Sigma_{t|t-1}). \tag{41}$$

In eq. (41), $A \in \mathbb{R}^{d\times d}$ denotes the *state transition matrix*, which maps the state of X_{t-1} onto the conditional expectation of X_t given by $\mu_{t|t-1} := AX_t$. Further, $\Sigma_{t|t-1} \in \mathbb{R}^{d\times d}$ denotes the *state disturbance matrix*, which, embedding the assumption of conditional independence of the state random variables (cf.eq. (4)), takes spherical form

$$\Sigma_{t|t-1} := \sigma^2_{t|t-1} I_d,\ \sigma^2_{t|t-1} > 0. \tag{42}$$

Note that in their structural formulation,

$$X_t = AX_{t-1} + \varepsilon_t \text{ with } p(\varepsilon_t) = N(\varepsilon_t; 0, \Sigma_{t|t-1}), \tag{43}$$

LGSs are sometimes also referred to as autoregressive processes (e.g., Tegmark, 2016)).

Effect and cause repertoires

For an LGS, the effect repertoire is given by

$$p_e(X_t|X_{t-1}) = N(X_t; AX_{t-1}, \Sigma_{t|t-1}) \tag{44}$$

and thus fully specified in terms of the system's transition and disturbance matrices. Because the concept of a uniform distribution is not meaningful in the context of continuous state variables, the specification of the cause repertoire in an LGS is more intricate. A standard option from the theory of Bayesian statistics is to use a zero-centred Gaussian with high variance for the uncertain distribution, i.e. to set

$$p_u(X_{t-1}) := N(X_{t-1}; 0, \Sigma_{t-1}) \tag{45}$$

with $\Sigma_{t-1} := \sigma_t^2 I_d$ and large σ_t^2 (Broemeling, 1984; Gelman et al., 2014). Based on the theory of multivariate Gaussian distributions, the cause repertoire under this assumption evaluates to

$$p_c(X_{t-1}|X_t) = N(X_{t-1}; \mu_{t-1|t}, \Sigma_{t-1|t}), \tag{46}$$

where

$$\mu_{t-1|t} = \Sigma_{t-1|t}\left(A^T \Sigma_{t|t-1} X_t\right) \text{ and } \Sigma_{t-1|t} = \left(\Sigma_t^{-1} + A^T \Sigma_{t|t-1}^{-1} A\right)^{-1}. \tag{47}$$

For a summary of these results from the theory of multivariate Gaussian distributions, please refer to Appendix C. In the limit of $\sigma_t^2 \to \infty$, and thus maximal uncertainty for the distribution $p_u(X_{t-1})$, eq. (47) simplifies to

$$\mu_{t-1|t} = \Sigma_{t-1|t}\left(A^T \Sigma_{t|t-1} X_t\right) \text{ and } \Sigma_{t-1|t} = \left(A^T \Sigma_{t|t-1}^{-1} A\right)^{-1} \tag{48}$$

In other words, one may disregard Σ_t^{-1} in eq. (47), because for very large σ_t^2, the precision Σ_t^{-1} approaches zero (e.g., Friston et al., 2002)). Note, however, that depending on the form of A, the covariance matrix $\Sigma_{t-1|t}$ in (47) may be undefined, resulting in an undefined value of ϕ. One of these cases is an all-zero matrix A, which is intuitive because there are no possible system transitions at all (and thus one might say that "there is no system to begin with" (Oizumi et al., 2014).

System decomposition

To compute ϕ for a system state $X \in \mathbb{R}$, the system under consideration has to be decomposed as discussed in Section 2 *Theory*. For an LGS, the crucial model component that induces stochastic dependencies between system state vectors X_{t-1} and X_t is the system transition matrix A. In the following, we discuss how the system transition matrix has to be manipulated in order to construct the set of factorized probability distributions $p_{ce}^{(i)}(X_{t-1}, X_t)$ (cf. eq. (16)). As described above, using the standard theory of multivariate Gaussians, this set of distributions can then directly be employed to compute the decomposed effect and cause repertoires.

We first note that for Gaussian random variables, uncorrelatedness (zero covariance) is equivalent to stochastic independence (Lauritzen, 1996). For LGSs, the covariances between the constituents of the random vectors X_{t-1} and X_t are encoded in the off-diagonal block matrices of the vectors' joint distribution covariance matrix (cf. Appendix C), i.e.

$$\mathrm{cov}(X_{t-1}, X_t) = A\Sigma_{t-1} \in \mathbb{R}^{d \times d}. \tag{49}$$

Setting elements in the covariance matrix $\mathrm{cov}(X_{t-1}, X_t)$ to zero thus corresponds to factorizing the joint distribution $p_{ce}(X_{t-1}, X_t)$. However, to compute ϕ, a well-defined mapping from the desired factorization (corresponding to a bipartition of the set (X_{t-1}, X_t)) to the form of the ensuing decomposed transition matrix $A^{(i)}$ is required. We next discuss a constructive scheme that achieves this.

Let $n := 2d$. We first note that by enumerating each random variable in (X_{t-1}, X_t) from 1 to n, each bipartition $\Pi^{(i)} = \left(\Pi_1^{(i)}, \Pi_2^{(i)}\right)$ of (X_{t-1}, X_t) can equivalently be represented by a set of two index sets

$$I^{(i)} := \left(I_1^{(i)}, I_2^{(i)}\right), i = 1, ..., 2^{n-1} - 1, \tag{50}$$

where $I_1^{(i)}, I_2^{(i)} \subset \mathbb{N}_n$ and $I_1^{(i)}$ and $I_2^{(i)}$ comprise the indices of the elements in $\Pi_1^{(i)}$ and $\Pi_2^{(i)}$, respectively. With the properties of bipartitions, it then follows immediately that

$$I_1^{(i)} \neq \emptyset, I_2^{(i)} \neq \emptyset, I_1^{(i)} \cap I_2^{(i)} = \emptyset, \text{ and } I_1^{(i)} \cup I_2^{(i)} = \mathbb{N}_n. \tag{51}$$

Let $\mathcal{I}_n$ denote the set of all possible bipartition index sets $I^{(i)}$. Using this index set representation of the bipartitions defined in eq. (13), we then define the mapping from bipartitions to system transition matrix variants by setting

$$f : \mathcal{I}_n \to \mathbb{R}^{n \times n}, \Pi^{(i)} \mapsto f(\Pi^{(i)}) := \mathrm{cov}(X_{t-1}, X_t)^{(i)}, \tag{52}$$

where with

$$A^{(i)} := v_{I_1^{(i)}} v_{I_1^{(i)}}^T + \left(1_d - v_{I_1^{(i)}}\right)\left(1_d - v_{I_2^{(i)}}\right)^T \tag{53}$$

and $v_{I_m^{(i)}} \in \mathbb{R}^d$ for $m = 1, 2$ given by

$$v_{I_m^{(i)}} := \left(v_{I_m^{(i)}}\right)_{1 \leq j \leq d} := \begin{cases} 1 & , j \in I_m^{(i)} \\ 0 & , j \notin I_m^{(i)}, \end{cases} \tag{54}$$

we instantiate the decomposed system covariance matrix by setting

$$\mathrm{cov}(X_t, X_{t-1})^{(i)} := A^{(i)} \circ \Sigma_{t-1}, \tag{55}$$

$\circ$ denoting the Hadamard matrix product. Note that in the case of an undefined Σ_t, the necessary system decomposition is still well-defined in terms of the parameters of the cause repertoire (cf. eq. (47)). To illustrate the mapping of eqs. (52) to (55), we visualize all ensuing covariance structures $\mathrm{cov}(X_{t-1}, X_t)^{(i)}$ for each bipartition of (X_{t-1}, X_t) in a three-dimensional LGS in Figure S5. Note that the components of X_{t-1} are indexed by $1, 2, 3$ and the components of X_t are indexed by $4, 5, 6$ (cf. eq. (50)). Black tiles indicate zero elements and white tiles indicate arbitrary elements of $\mathrm{cov}(X_{t-1}, X_t)^{(i)}$.

Computation of ϕ

Based on the above, we are now in the position to evaluate ϕ for exemplary LGSs. We here limit ourselves to some exploratory simulations in the class of two-dimensional (planar) LGSs (e.g., Meiss, 2007; Hirsch et al., 2012; Arnold, 2013). While the exhaustive evaluation of ϕ over state space is no longer feasible, we can compute ϕ for exemplary state trajectories. To this end, we consider planar LGSs and study how exemplary parameter changes and their associated state trajectories variations qualitatively affect ϕ. The simulations are visualized in Figure 5. Each of the six subpanels A - F instantiates a different parameter setting. Within each subpanel, the left panel visualizes the system evolution in

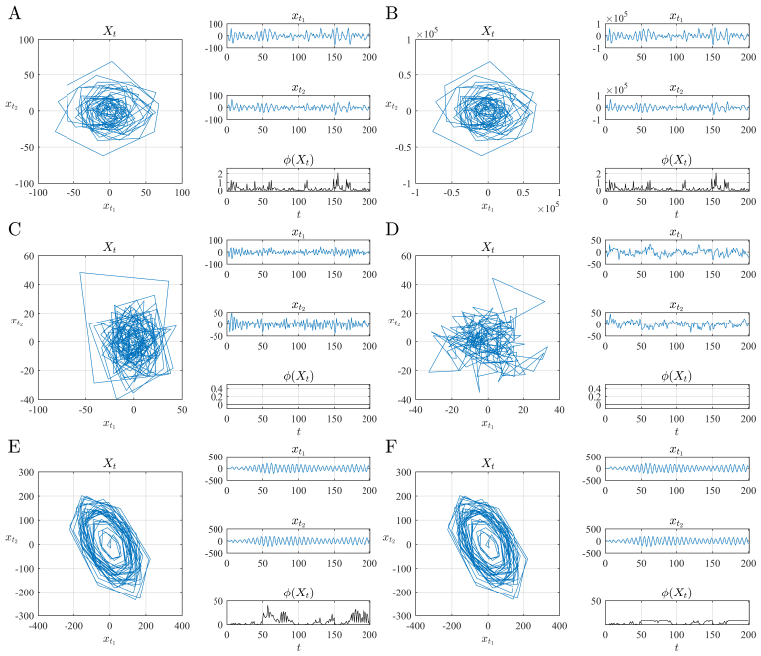

Figure 5: Evaluation of ϕ in exemplary linear Gaussian systems. Within each figure panel A-F, the left panel visualizes the system evolution in its two-dimensional state space, while the right side displays the trajectory of each system component and the evolution of ϕ. Panel A shows an exemplary LGS, defined by the parameters of (56). Increasing the variance in the disturbance matrix by some orders of magnitude results in a scaling effect of the system trajectories (B), without a considerable effect on ϕ. In contrast, altering the transition matrix has pronounced effects on ϕ. Setting the diagonal (C) or off-diagonal (D) elements of A in(56) to zero results in markedly different state trajectories and the vanishing of $\phi = 0$ for all observed states. Panels E and F demonstrate that the system parameters may also interact. Here, the state transition matrix of (57) increases the sensitivity of ϕ to alterations of the disturbance and precision matrix, as shown in panel F, where precision is decreased by four orders of magnitude.

$\mathbb{R}^2$ for $t = 1, ..., 200$ time steps. The upper right two panels visualize the axes projections of the system state, reminiscent of empirical data recording channels (e.g. EEG electrodes or fMRI voxel time-courses). The lowermost panel depicts ϕ as a function of the temporal evolution of the system state.

Panel A of Figure 5 visualizes the state evolution of a system with state transition, state disturbance, and marginal precision matrices

$$A := \begin{pmatrix} 0.5 & -1 \\ 0.5 & 0.5 \end{pmatrix}, \Sigma_{t|t-1} := \begin{pmatrix} 10^2 & 0 \\ 0 & 10^2 \end{pmatrix}, \text{ and } \Sigma_t^{-1} := \begin{pmatrix} 10^{-1} & 0 \\ 0 & 10^{-2} \end{pmatrix}, \quad (56)$$

respectively with initial state $X_0 := (1, 1)^T$. Changing the state disturbance matrix by six orders of magnitude to $\Sigma_{t|t-1} := 10^8 I_2$ only affects the scale of the trajectory, but does not affect ϕ, as shown in Figure 5B. Likewise, for this system, different choices of the marginal precision matrix Σ_t leave ϕ unchanged (data not shown). Naturally, changing the state transition matrix has profound effects on both the system state trajectory and the observed values of ϕ. In Figure 5C and D, a system with parameters as in (56) is depicted, where the diagonal and off-diagonal entries of the state transition matrix A were set to zero, resulting in the vanishing of ϕ over the simulation interval. Finally, the state transition and marginal precision parameters may also interact in their effect on ϕ: Panels E and F of Figure 5 visualize simulations of a planar system with state and state disturbance matrices

$$A := \begin{pmatrix} 1 & 1 \\ -1 & 1 \end{pmatrix} \text{ and } \Sigma_{t|t-1} := \begin{pmatrix} 10^2 & 0 \\ 0 & 10^2 \end{pmatrix}. \quad (57)$$

Changing the marginal precision matrix from $\Sigma_t^{-1} := 10^{-2}$ (Figure 5 E) to $\Sigma_t^{-1} := 10^{-6}$ (Figure 5F) in this system has quantitative effects on ϕ, reducing its peaks to a plateau-like behaviour, while retaining its qualitiative properties.

In summary, based on the above formulation, ϕ can be evaluated in LGSs, as demonstrated by simulations, shows a rich temporal diversity, and exhibits a dependence on the LGS parameter settings. In the further development of IIT, the class of planar LGSs may form a useful analytical test bed for the qualitative and quantitative dependence of ϕ on a system's parameterization, because, at least in the deterministic scenario, planar dynamical systems are well understood (e.g., Meiss, 2007; Hirsch et al., 2012).

4 Discussion

In the present work, we have developed a general probabilistic formulation of integrated information ϕ, starting from the most recent instantiation of IIT (Oizumi et al., 2014). All necessary mathematical operations in the derivation of ϕ are based on a system's joint distribution $p_{ce}(X_{t-1}, X_t)$ over two adjacent points in time. We presented a constructive rule for the system decomposition into two subsystems, which corresponds to a flexible factorization of this joint distribution. Our formulation is readily applied to non-parametric discrete state systems, as validated in the exemplary system from IIT 3.0, in which we also show that ϕ mirrors state evolution over time. As a consequence of our formulation, undefined conditional distributions are eschewed and virtual elements are not necessary. Furthermore, to make progress towards the computation of

ϕ based on experimental data, we applied our formulation to continuous variables in the versatile model class of LGSs. For the Gaussian case, we approach the problem of a maximum uncertainty distribution over past states by defining a prior distribution with near-zero precision, in line with standard Bayesian strategies. We showed that, for LGSs, the system decomposition corresponds to mapping all possible partitions to covariance matrix forms of system states at adjacent time points. Finally, we explored the qualitative effect of different LGS parameters on integrated information and evaluate the evolution of system states and ϕ over time, revealing non-trivial temporal dynamics. In summary, our formulation is - at least in principle - general enough to be applicable to different model classes and investigate the dynamics of these models. In the following, we start by discussing several related studies and then turn to some open questions regarding the computation of integrated information.

The theoretical development of IIT has so far mainly rested on discrete state systems (Tononi, 2004, 2008; Balduzzi and Tononi, 2008, 2009; Tononi, 2012; Oizumi et al., 2014). However, this model class is not suitable in most empirical settings, because neuroimaging and other functional data are usually not limited to discrete states. In consequence, a number of efforts have recently been targeted at the evaluation of ϕ in for continuous state (Barrett and Seth, 2011; Oizumi et al., 2016; Mediano et al., 2016; Tegmark, 2016).

Barrett and Seth (2011) were among the first in this regard and propose a measure of integrated information that does allow for the evaluation of integration in time-series data. The main difference to the formulation presented here is that Barrett and Seth focus on an earlier version of the theory - integrated information theory 2.0 (Balduzzi and Tononi, 2008) - which is based on entropy-related measures, while we start explicitly from the most recent instantiation of IIT. As has been pointed out by Tegmark (2016) and Oizumi et al. (2016), however, the integrated information measure proposed by Barrett and Seth (2011) can become negative. This somewhat complicates its interpretation, although the interesting question has been raised of whether "negative integration" could reflect redundancy in a system (Barrett, 2015). Note that ϕ as presented herein is well-bounded by zero because both the EMD and the KL divergence cannot be negative (Levina and Bickel, 2001; Cover and Thomas, 2012). Introducing theoretical requirements for the lower and upper bounds for ϕ was one of the main motivations for recent work by Oizumi et al. (2016), who also evaluate their measure of information integration in monkey electrocorticogram data. The major differences to our work are that the authors develop their measure based on IIT 2.0 and involves "atomic partitioning", which we discuss in more detail below.

An impressive collection of different mathematical options to measure information integration was presented recently by Tegmark (2016). Similarities between our approach and that of Tegmark (2016) include the application of ϕ to Markov process system models and and the ensuing construction of cause and effect repertoires. We also use the Kullback-Leibler Divergence for the Gaussian case, because the EMD is computationally infeasible for continuous variables. We differ, however, in the constructive algorithm for the system decomposition and in that we do define ϕ as presented in IIT 3.0 for the continuous case by using a (near-)zero precision distribution over past states. Importantly, we consider all possible partitions in the system decomposition, both symmetrical and asymmetrical, which we discuss in more detail below.

In summary, while there is recent interest in the evaluation of integrated information based on experimental data, so far there has been no convergence on a gold-standard measure. This is of course closely related to the fact that the theoretical framework of integrated information itself, and specifically in its application to conciousness research, remains under development. For the remainder of the discussion we focus on three aspects that in our opinion leave room for further consolidation of the theory.

The normalization problem

Some concerns have been raised regarding the necessity of normalization in IIT. First of all, it is useful to distinguish what needs to be normalized. The Scholarpedia entry for IIT introduces a normalization constant for the cause repertoires, i.e. $p_c(X_{t-1}|X_t)$, in order to ensure that "the repertoire sums to one" (Tononi, 2015). A similar normalization constant is not introduced for the effect repertoire $p_e(X_t|X_{t-1})$. However, it is the very nature of conditional probability distributions that they sum to one, so the question is why normalization for $p_c(X_{t-1}|X_t)$ is required in the first place. In our formulation, normalization is not necessary to ensure unity. However, normalization may become necessary if one introduces conditional distributions over the empty set of variables, as detailed above. Essentially, this corresponds to the case distinction of whether the marginal distribution for conditioning factorizes or not (cf. eq. (21)). Whether this is the case in turn depends on the particular system decomposition, i.e. whether or not there are variables for conditioning in both subsets $\Pi_1^{(i)}$ and $\Pi_1^{(i)}$ according to bipartition i.

A second question is whether ϕ itself needs to be normalized. The rationale behind this is that ϕ is thought to get small whenever one of the subsets in a bipartition is small. Since ϕ is defined over the factorization that makes the least difference compared to the unfactorized system (the minimum information partition), ϕ may be biased to asymmetrical partitions. Indeed, in previous formulations of IIT, this issue was explicitly addressed by introducing a normalization factor favouring balanced partitions (Balduzzi and Tononi, 2008), while IIT 3.0 does not account for this. In fact, Tononi states in his response to Aaronson that "normalization is not necessary in IIT 3.0" (Tononi, 2014). However, it remains unclear why this should be the case. Since the normalization issue is tightly linked to the partition problem, we will treat the latter in detail below.

The partition problem

The partition problem is of interest both on theoretical grounds and in the attempt to make the calculation of ϕ feasible for experimental data. As shown in Appendix B, the number of unique bipartitions of any set with cardinality n is given by an identity of the Stirling number of the 2^{nd} kind. The problem is that this number grows exponentially with n, thus seriously compromising the computational tractability for sets containing more than about $15 - 20$ dimensions. There have been several approaches to this problem. Tegmark (2016) limits himself to symmetrical bipartitions. While this has the advantage of excluding uneven partitions by definition, thus circumventing the issue of normalization, most partitions in a large set of elements will be symmetrical (e.g., Figure 6).

This approach thus concedes generality but does not gain much computational tractability. It is furthermore unclear how this will fare in networks containing an odd number of elements as these have no strictly symmetrical bipartitions. In the study discussed above, Oizumi et al. (2016) have used atomic partitions, i.e., the complete factorization of conditional distributions over all variables, to calculate a modified measure of integrated information in primate electrocorticogram data. The authors concede that atomic partitioning and to a lesser extent also symmetrical partitioning will overestimate information integration because it tends to maximize rather than minimize the informational difference the partition makes to the system. Thus, the question is which partitioning approach is best when the theoretical analysis of all bipartitions is no longer feasible. Our formulation does take all theoretically possible partitions into account and is applicable to continuous data, but suffers from the same limitations of combinatorial explosion as previous endeavours. While we do not yet offer a definite remedy to this problem, there are a couple of outlooks that may be worth discussing:

First, consider Figure 6, in which we depict the subset sizes of all bipartitions of a set with cardinality 18. The bulk of the $k = 131071$ possible bipartitions comprises 24310 symmetrical bipartitions (i.e. $|\Pi_{1/2}| = 9$), while there are only 18 maximally asymmetrical bipartitions ($|\Pi_{1/2}| = 1$) and 153 submaximally asymmetrical bipartitions ($|\Pi_{1/2}| = 2$). If one is willing to accept the current definition of ϕ as the (unnormalized) minimum over all bipartitions, we may make a virtue of necessity and limit ourselves not to symmetrical bipartitions but rather to more asymmetrical bipartitions, as they will tend to result in the minimal information partition. We could simply threshold the cardinality of the subsets we wish to consider, thus governing how many of the more symmetrical partitions to take into account as well. Note that maximally asymmetrical partitions also reduce computations to linear time. As opposed to atomic partitions, they are bipartitions, however, and will tend to under- rather than overestimate integration. However, such an approach raises the question to what extent the definition of ϕ then provides additional information compared to information that is contained in a system's transition probability distribution. As an example, consider the definition of ϕ_e. ϕ_e is defined as the minimum divergence between the distribution $p_c(X_t|X_{t-1} = X)$ of the original system and the distributions $p_c^{(i)}(X_t|X_{t-1} = X)$ upon removal of possible stochastic dependencies of the original system. Naturally, this divergence becomes minimal if these distributions are most similar, which in turn is the case if the transition probability distributions of the original system and the transition probability distribution of the decomposed system variant are most similar. In other words, the minimization procedure primarily identifies the transition probability distribution of the original system. However, this was assumed to be known from the beginning. It may be thus more straight-forward to define a measure of information integration directly on the (estimated) transition probability distribution of a system, rather than having to evaluate all possible decompositions of a system only to identify the original system transition probability distribution. This could also alleviate the problem of having to define a maximum uncertainty distribution over past states in the Gaussian case, for which we could simply take the marginal distribution. Note that a similar observation regarding the minimization procedure has been raised by Tegmark (2016).

A second approach is to use our formulation to start evaluating ϕ on a macroscopic scale, i.e. over merely a few brain regions of interest containing lots of neural elements, thus circumventing combinatorial explosion by scaling. Indeed, such large-scale "hot zones" for the neural correlates of consciousness have been identified in posterior cortical zones in recent years (e.g. Koch et al., 2016). While this disregards the recently developed concept of causal emergence on different spatio-temporal scales (Hoel et al., 2013; Tononi et al., 2016), it is a start, and similar approaches are successfully used in dynamic causal modelling (Friston et al., 2003) as well as graph-theory and mean-field-based measures (Deco et al., 2015) that try to capture the behaviour of large-scale brain networks.

A third approach could be to find an estimate of which partitions are likely to result in a great difference to the original system and then discard these. For example, in a complex system, one could analyse the network structure based on graph theoretical measures and then reduce the number of partitions that have to be evaluated by discarding all partitions that cut through the connections (i.e. introduce stochastic independence) of a hub node. Since hub nodes are by definition strongly connected and can thus be assumed to "make a difference to the system" (Oizumi et al., 2014), such an approach could both substantially reduce the computational load (because many partitions would affect the hub) and serve as a theoretically and biologically plausible approximation.

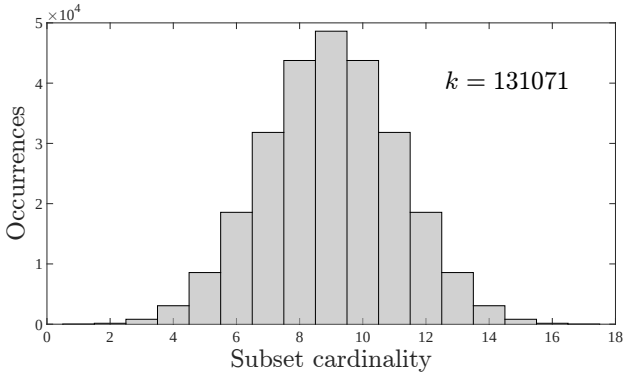

Figure 6: Frequency histogram of the subset sizes of all possible bipartitions of a set with cardinality 18. The distribution is completely symmetrical (each subset with cardinality 1 corresponds to another with cardinality 17 as they belong to the same partition, and so on until $|\Pi_{1/2}| = 9$). Symmetrical bipartitions make up the majority of all possible partitions while maximally asymmetrical partitions grow linearly with n.

State-dependency and temporal evolution of ϕ

Finally, we consider the state-dependency and the temporal dynamics of ϕ. Recent work by Virmani and Nagaraj (2016) aims at bridging the gap between IIT and the perturbational complexity index, a practical measure based on TMS stimulation with simultaneous EEG recordings that has been successfully applied in the clinical quantification of consciousness (Casali et al., 2013). To this

end, the authors derive a measure of compression-complexity which is calculated by a maximum entropy perturbation of each node in an atomic partition. While this is certainly a promising approach, their measure shows minimal dependency on the current system state, similar to the measure proposed by Barrett and Seth (2011). Based on empirical studies (Koch et al., 2016), however, it seems likely that any measure that quantifies consciousness should indeed be state-dependent (cf. (Tegmark, 2016)). On theoretical grounds, state-independence would also violate the selectivity postulate that IIT propoposes as a prerequisite for information (e.g. Figure 3 in (Oizumi et al., 2014)). Moreover, insensitivity to system states raises the question of whether a measure in fact represents integration or rather the capacity to integrate (Barrett and Seth, 2011). In the present work, state-dependency is preserved because the conditional distributions over X_{t-1} and X_{t+1} are always evaluated for $X_t = X$ (cf. eqs. (22) and (23)). Furthermore, since it is the state at time t that defines the cause and effect repertoire, the decomposition of which in turns defines ϕ, it is not surprising that ϕ essentially mirrors the state evolution in simple discrete state systems like the one evaluated herein. In the exemplary linear Gaussian system, the temporal dynamics of ϕ are more complex. Here, ϕ shows a rich diversity over the oscillatory state trajectories of the system components, and it will be interesting to investigate under which circumstances ϕ assumes high values in such systems. In this regard, a recent study by Mediano et al. (2016) and colleagues has related integrated information to metastability in a system of coupled Kuramoto oscillators (Kuramoto, 2012). Investigating the temporal dynamics of ϕ in such systems appears to us as a promising approach, since metastability has been suggested as a potential mechanistic explanation for the emergence of large-scale functional brain networks (Deco et al., 2015) and because it offers a potential link to the dynamic causal modelling framework which has long been extended to the analysis of phase-coupled data (Penny et al., 2009).

Conclusion

To conclude, with the general probabilistic formulation of ϕ presented herein, we hope to make a contribution to the traceability, parsimony, and generality of integrated information theory. Our formulation is readily applicable to discrete state systems and generalizes the application of integrated information theory to different probabilistic model classes, such as linear Gaussian systems - thereby enhancing the potential transfer of integrated information theory to the context of functional neuroimaging data analysis. We thus consider our approach as a step towards the evaluation of integrated information from experimental data with the prospect of investigating its temporal dynamics and as a basis for the further development of the theory of integrated information in an explicit probabilistic framework.

Appendices

Appendix A

The re-conceptualization of integrated information theory as "integrated information theory 3.0" by Oizumi et al. (2014) has introduced various nuances and refinements to measuring integrated information. This has resulted in a variety of "phi" measures, which we summarize here and relate to the measure ϕ as discussed in the main text.

In Oizumi et al. (2014), φ denotes the minimum earth movers distance (EMD) (Levina and Bickel, 2001; Mallows, 1972), over all possible bipartitions of a system, referred to as "minimum information partition" (MIP). It is thus also denoted by φ^{MIP}. Since we are exclusively interested in the MIP, we omit the superscript henceforth. In the temporal direction past, φ corresponds to all factorizations of the "cause repertoire" and thus reflects "integrated cause information", denoted by φ_{cause}. Similarly, the minimum EMD over all factorizations of the "effect repertoire" represents integrated effect information, denoted by φ_{effect}. The minimum of φ_{cause} and φ_{effect} is denoted by $\varphi_{cause-effect}$ in Oizumi et al. (2014) and corresponds to the measure we refer to as ϕ. Note that we chose this variant notation for unique reference to our formulation.

In addition to ϕ, Oizumi et al. (2014) introduce a number of additional "phi" measures, which we list here for completeness. Firstly, φ^{max}_{cause} and φ^{max}_{effect} are the respective maxima over all φ_{cause} and φ_{effect}, which are assessed over the powerset of all possible conditional distributions of a set of elements. See for example Figure 8 in Oizumi et al. (2014). Together, these define φ^{max} as the minimum of φ^{max}_{cause} and φ^{max}_{effect}. Secondly, Φ (or more precisely Φ^{MIP}) is a weighted sum of all EMDs corresponding to φ^{max}_{cause} and φ^{max}_{effect} compared to the "unconstrained past and future repertoires". In Oizumi et al. (2014), this measure is also referred to as "integrated conceptual information". Finally, Φ^{max} is the maximum of all Φ^{MIP} over the powerset of all original sets to be considered. In a network containing n nodes, the greatest value for Φ^{MIP} may thus be found over a subset with cardinality $c < n$.

Appendix B

Finding the number of partitions of a set containing a finite number of elements is a well-known problem in combinatorics (Cameron, 1994). The number of all bipartitions of a finite set can be evaluated by the Stirling number of the 2^{nd} kind,

$$S_{n,m} = \frac{1}{m!} \sum_{i=0}^{m} (-1)^{m-i} \binom{m}{i} i^n, \tag{58}$$

which counts the number of possible partitions of a set of cardinality $n \in \mathbb{N}$ into $m \in \mathbb{N}$ non-empty disjoint subsets. For the purposes of IIT, bipartitions are

required. Thus, $m = 2$ and we have

$$\begin{aligned} S(n,2) &= \frac{1}{2!}\sum_{i=0}^{2}(-1)^{2-i}\binom{2}{i}i^n \\ &= \frac{1}{2!}\left(\left((-1)^{2-0}\binom{2}{0}0^n\right)+\left((-1)^{2-1}\binom{2}{1}1^n\right)+\left((-1)^{2-2}\binom{2}{2}2^n\right)\right) \\ &= \frac{1}{2}\left(\left((-1)^2(1)\,0^n\right)+\left((-1)^1(2)\,1^n\right)+\left((-1)^0(1)\,2^n\right)\right) \\ &= \frac{1}{2}\left(0+(-2)+2^n\right) \\ &= -1+\frac{2^n}{2} \\ &= 2^{n-1}-1 \end{aligned} \tag{59}$$

In the main text, we use the denotation $k := 2^{2d-1} - 1$ for a set of cardinality $n = 2d$.

Appendix C

In this Appendix, we collect a number of well-known results from the theory of multivariate Gaussian distributions. We state these results in theorem form. For proofs, see (Bishop, 2006; Barber, 2012) and (Murphy, 2012).

Theorem: Gaussian joint, conditional, and marginal distributions.

Given a Gaussian marginal density

$$p(z) = N(z;\mu_z,\Sigma_z), \text{where } z,\mu_z \in \mathbb{R}^p, \Sigma_z \in \mathbb{R}^{p\times z} \tag{60}$$

and a Gaussian conditional distribution density

$$p(y|z) = N(y;A\mu_z,\Sigma_{y|z}), \text{where } y \in \mathbb{R}^n, A \in \mathbb{R}^{n\times p}, \Sigma_{y|z} \in \mathbb{R}^{n\times n} \tag{61}$$

the joint density of y and z

$$p(y,z) = p(y|z)p(z) \tag{62}$$

is given by

$$p(y,z) = N\left(\begin{pmatrix}y\\z\end{pmatrix};\mu_{y,z},\Sigma_{y,z}\right), \text{ with } (y,z)^T, \mu_{y,z} \in \mathbb{R}^{n+p}, \Sigma_{y,z} \in \mathbb{R}^{n+p\times n+p} \tag{63}$$

where

$$\mu_{y,z} = \begin{pmatrix}A\mu_z\\ \mu_z\end{pmatrix} \text{ and } \Sigma_{y,z} = \begin{pmatrix}\Sigma_{y|z}+A\Sigma_z A^T & A\Sigma_z\\ \Sigma_z A^T & \Sigma_z\end{pmatrix} \tag{64}$$

Further, the conditional distribution

$$p(z|y) = \frac{p(y,z)}{p(y)} \tag{65}$$

is given by the Gaussian probability density

$$p(z|y) = N(z;\mu_{z|y},\Sigma_{z|y}), \tag{66}$$

where

$$\mu_{z|y} = \Sigma_{z|y}(\Sigma_z^{-1}\mu_z + A^T\Sigma_{y|z}^{-1}y) \in \mathbb{R}^p \text{ and } \Sigma_{z|y} = (\Sigma_z^{-1} + A^T\Sigma_{y|z}^{-1}A)^{-1} \in \mathbb{R}^{p\times p}, \tag{67}$$

and the marginal distribution

$$p(y) = \int p(y,z)dz \tag{68}$$

is given by the Gaussian probability density

$$p(y) = N(y; \mu_y, \Sigma_y), \tag{69}$$

where

$$\mu_y = Az \in \mathbb{R}^n \text{ and } \Sigma_y = \Sigma_{y|z} + A\Sigma_z A^T \in \mathbb{R}^{n\times n}. \tag{70}$$

□

References

Aaronson, S. (2014). Why i am not an integrated information theorist (or, the unconscious expander).

Arnold, L. (2013). *Random dynamical systems*. Springer Science & Business Media.

Balduzzi, D. and Tononi, G. (2008). Integrated information in discrete dynamical systems: Motivation and theoretical framework. *PLoS Comput Biol*, 4(6):e1000091.

Balduzzi, D. and Tononi, G. (2009). Qualia: the geometry of integrated information. *PLoS Comput Biol*, 5(8):e1000462.

Barber, D. (2012). *Bayesian Reasoning and Machine Learning*. Cambridge University Press.

Barrett, A. B. (2015). Exploration of synergistic and redundant information sharing in static and dynamical gaussian systems. *Phys. Rev. E*, 91:052802.

Barrett, A. B. and Seth, A. K. (2011). Practical measures of integrated information for time-series data. *PLoS Comput Biol*, 7(1):e1001052.

Billingsley, P. (2008). *Probability and measure*. John Wiley & Sons.

Bishop, C. M. (2006). *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA.

Broemeling, L. D. (1984). *Bayesian Analysis of Linear Models*. Statistics: A Series of Textbooks and Monographs. Taylor & Francis.

Cameron, P. J. (1994). *Combinatorics: topics, techniques, algorithms*. Cambridge University Press.

Casali, A. G., Gosseries, O., Rosanova, M., Boly, M., Sarasso, S., Casali, K. R., Casarotto, S., Bruno, M.-A., Laureys, S., Tononi, G., and Massimini, M. (2013). A theoretically based index of consciousness independent of sensory processing and behavior. *Science Translational Medicine*, 5(198):198ra105–198ra105.

Cerullo, M. A. (2015). The problem with phi: A critique of integrated information theory. *PLoS Comput Biol*, 11(9):e1004286.

Cover, T. M. and Thomas, J. A. (2012). *Elements of information theory*. John Wiley & Sons.

Cox, D. and Miller, H. (1977). *The Theory of Stochastic Processes*. Science paperbacks. Taylor & Francis.

Dawid, A. P. (1979). Conditional independence in statistical theory. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 1–31.

Deco, G., Tononi, G., Boly, M., and Kringelbach, M. L. (2015). Rethinking segregation and integration: contributions of whole-brain modelling. *Nature Reviews Neuroscience*, 16(7):430–439.

Efron, B. and Hastie, T. (2016). *Computer Age Statistical Inference*, volume 5. Cambridge University Press.

Friston, K., Glaser, D., Henson, R. N. A., Kiebel, S., Phillips, C., and Ashburner, J. (2002). Classical and bayesian inference in neuroimaging: applications. *Neuroimage*, 16(2):484–512.

Friston, K., Harrison, L., and Penny, W. (2003). Dynamic causal modelling. *Neuroimage*, 19(4):1273–1302.

Geiger, D., Verma, T., and Pearl, J. (1990). Identifying independence in bayesian networks. *Networks*, 20(5):507–534.

Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B. (2014). *Bayesian data analysis*, volume 2. Chapman & Hall/CRC Boca Raton, FL, USA.

Hirsch, M. W., Smale, S., and Devaney, R. L. (2012). *Differential equations, dynamical systems, and an introduction to chaos*. Academic press.

Hoel, E. P., Albantakis, L., and Tononi, G. (2013). Quantifying causal emergence shows that macro can beat micro. *Proceedings of the National Academy of Sciences*, 110(49):19790–19795.

Jordan, M. I. (1998). *Learning in graphical models*, volume 89. Springer Science & Business Media.

Koch, C., Massimini, M., Boly, M., and Tononi, G. (2016). Neural correlates of consciousness: progress and problems. *Nat Rev Neurosci*, 17(5):307–321.

Kullback, S. and Leibler, R. A. (1951). On information and sufficiency. *The annals of mathematical statistics*, 22(1):79–86.

Kuramoto, Y. (2012). *Chemical oscillations, waves, and turbulence*, volume 19. Springer Science & Business Media.

Lauritzen, S. L. (1996). *Graphical models*, volume 17. Clarendon Press.

Levina, E. and Bickel, P. (2001). The earth mover's distance is the mallows distance: some insights from statistics. In *Proc. Eighth IEEE Int. Conf. Computer Vision ICCV 2001*, volume 2, pages 251–256 vol.2.

Mallows, C. (1972). A note on asymptotic joint normality. *The Annals of Mathematical Statistics*, pages 508–515.

Mediano, P. A. M., Farah, J. C., and Shanahan, M. (2016). Integrated information and metastability in systems of coupled oscillators. *arXiv preprint arXiv:1606.08313*.

Meiss, J. D. (2007). *Differential dynamical systems*, volume 14. Siam.

Murphy, K. P. (2012). *Machine learning: a probabilistic perspective*. MIT press.

Oizumi, M., Albantakis, L., and Tononi, G. (2014). From the phenomenology to the mechanisms of consciousness: Integrated information theory 3.0. *PLoS Comput Biol*, 10(5):e1003588.

Oizumi, M., Amari, S.-i., Yanagawa, T., Fujii, N., and Tsuchiya, N. (2016). Measuring integrated information from the decoding perspective. *PLoS Comput Biol*, 12(1):e1004654.

Ostwald, D., Kirilina, E., Starke, L., and Blankenburg, F. (2014). A tutorial on variational bayes for latent linear stochastic time-series models. *Journal of Mathematical Psychology*, 60:1–19.

Ostwald, D., Porcaro, C., and Bagshaw, A. P. (2010). An information theoretic approach to eeg-fmri integration of visually evoked responses. *Neuroimage*, 49(1):498–516.

Ostwald, D. and Starke, L. (2016). Probabilistic delay differential equation modeling of event-related potentials. *Neuroimage*, 136:227–257.

Penny, W., Litvak, V., Fuentemilla, L., Duzel, E., and Friston, K. (2009). Dynamic causal models for phase coupling. *J Neurosci Methods*, 183(1):19–30.

Roweis, S. and Ghahramani, Z. (1999). A unifying review of linear gaussian models. *Neural computation*, 11(2):305–345.

Tegmark, M. (2016). Improved measures of integrated information.

Tononi, G. (2004). An information integration theory of consciousness. *BMC Neurosci*, 5:42.

Tononi, G. (2005). Consciousness, information integration, and the brain. *Prog Brain Res*, 150:109–126.

Tononi, G. (2008). Consciousness as integrated information: a provisional manifesto. *The Biological Bulletin*, 215(3):216–242.

Tononi, G. (2012). Integrated information theory of consciousness: an updated account. *Arch Ital Biol*, 150(4):293–329.

Tononi, G. (2014). Why scott should stare at a blank wall and reconsider (or, the conscious grid).

Tononi, G. (2015). Integrated information theory. *Scholarpedia*, 10(1):4164.

Tononi, G., Boly, M., Massimini, M., and Koch, C. (2016). Integrated information theory: from consciousness to its physical substrate. *Nature Reviews Neuroscience*.

Virmani, M. and Nagaraj, N. (2016). A compression-complexity measure of integrated information. *arXiv preprint arXiv:1608.08450*.

Supporting Information

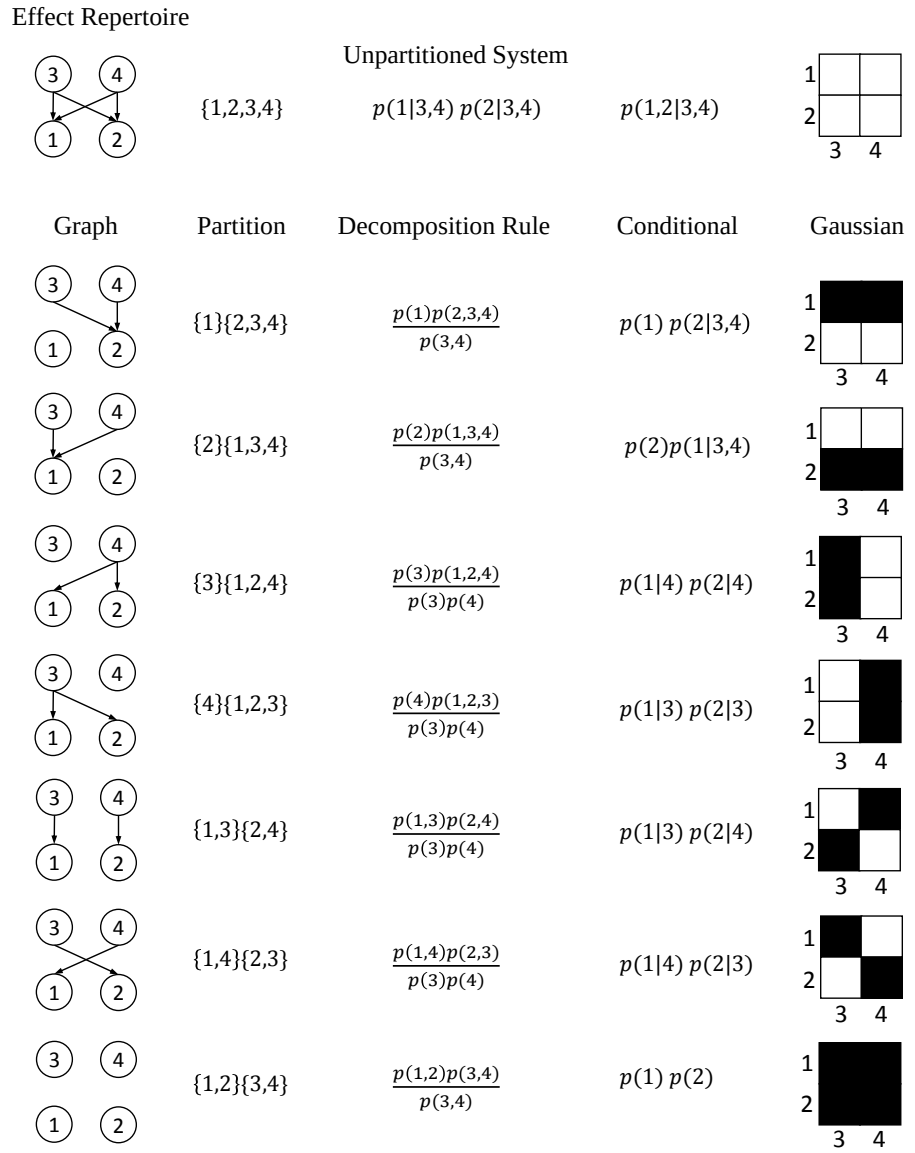

Figure S1: Exhaustive system decomposition (effect repertoire). Visualization of all $k = 2^{4-1} - 1 = 7$ unique bipartitions of a two-dimensional system with $X_t := \{1, 2\}$ and $X_{t-1} := \{3, 4\}$. For the effect repertoire (first row), the joint distribution over $\{1, 2, 3, 4\}$ is conditioned on $\{3, 4\}$. The graphical model for every partition is depicted in the left column alongside the bipartition of $\{1, 2, 3, 4\}$ in the remaining rows. The decomposition rule as presented in the main text is displayed along with the ensuing conditional distribution in the third and fourth columns. For the Gaussian case, the effects of the system decomposition on the covariance structure $\text{cov}(X_{t-1}, X_t)$ of the joint distribution are shown with black tiles indicating zero elements.

.

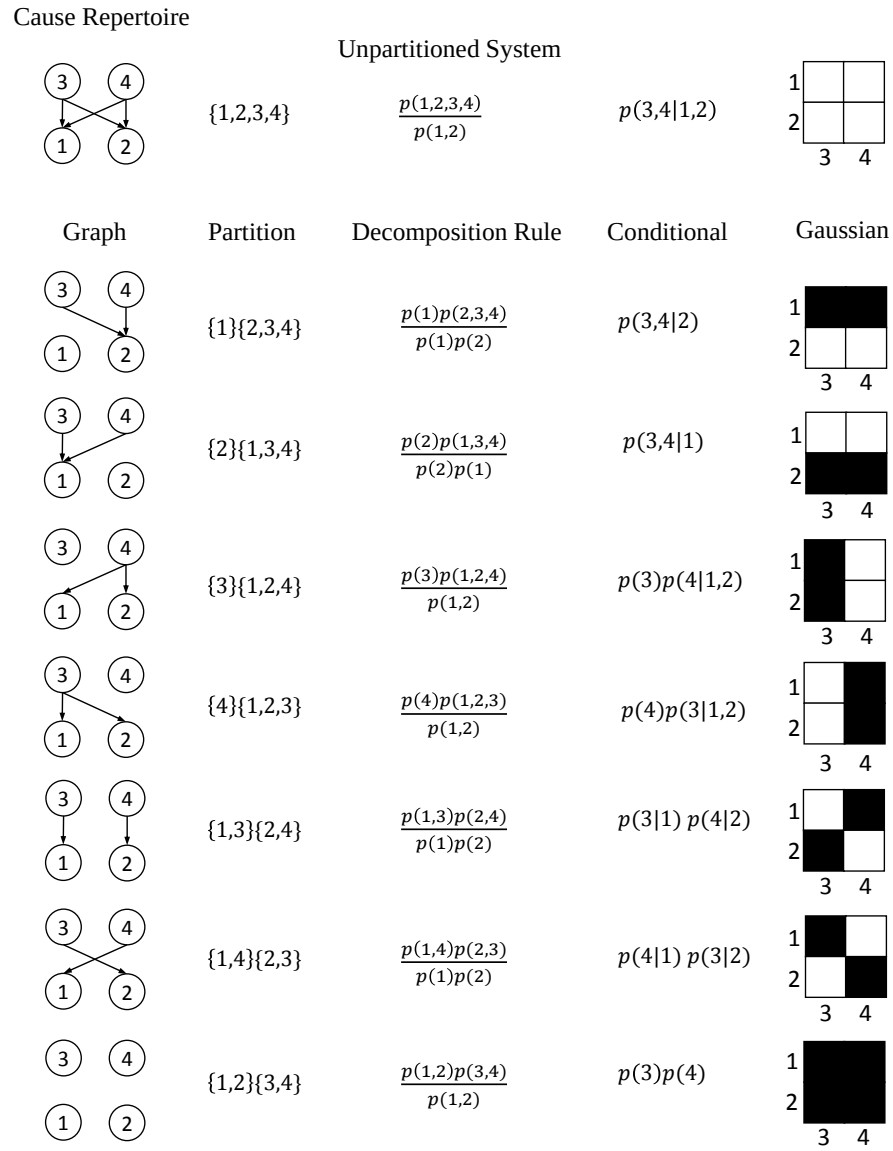

Figure S2: Exhaustive system decomposition (cause repertoire). Visualization of all $k = 2^{4-1} - 1 = 7$ unique bipartitions of a two-dimensional system with $X_t := \{1, 2\}$ and $X_{t-1} := \{3, 4\}$. For the cause repertoire (first row), the joint distribution over $\{1, 2, 3, 4\}$ is conditioned on $\{1, 2\}$. The graphical model for every partition is depicted in the left column alongside the bipartition of $\{1, 2, 3, 4\}$ in the remaining rows. The decomposition rule as presented in the main text is displayed along with the ensuing conditional distribution in the third and fourth columns. For the Gaussian case, the effects of the system decomposition on the covariance structure $\mathrm{cov}(X_{t-1}, X_t)$ of the joint distribution are shown with black tiles indicating zero elements.

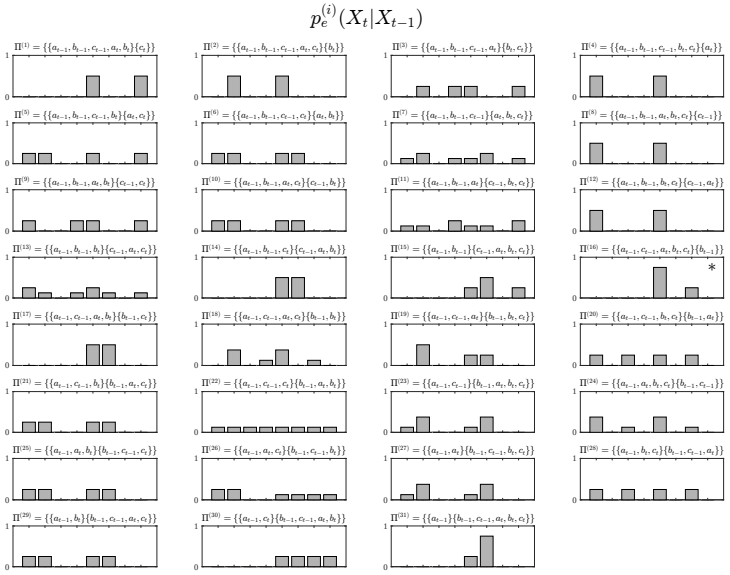

Figure S3: Decomposed effect repertoire for an exemplary discrete-state system. For the system state $X = (1, 0, 0)$, the figure visualizes the complete set of decomposed variants of the effect repertoire $p(X_t|X_{t-1} = X)$. Each subpanel includes the corresponding partition of the set of random variables $(X_{t-1}, X_t) = (a_{t-1}, b_{t-1}, c_{t-1}, a_t, b_t, c_t)$ that gives rise to these variants. The asterisk signifies the decomposed effect repertoire which results in the minimum EMD with respect to the original effect repertoire, which in turn corresponds the effect information $\phi_e((1, 0, 0))$.

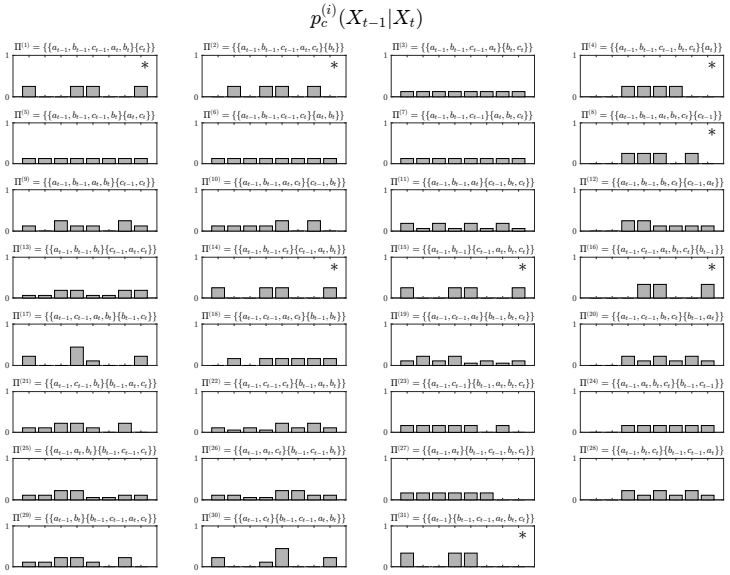

Figure S4: Decomposed cause repertoire for an exemplary discrete-state system. For the system state $X = (1, 0, 0)$, the figure visualizes the complete set of decomposed variants of the random variable set (X_{t-1}, X_t) giving rise to the cause repertoire variants. The asterisks signify the decomposed cause repertoires which results in a minimum EMD with respect to the original cause repertoire, which in turn corresponds the cause information $\phi_c((1, 0, 0))$

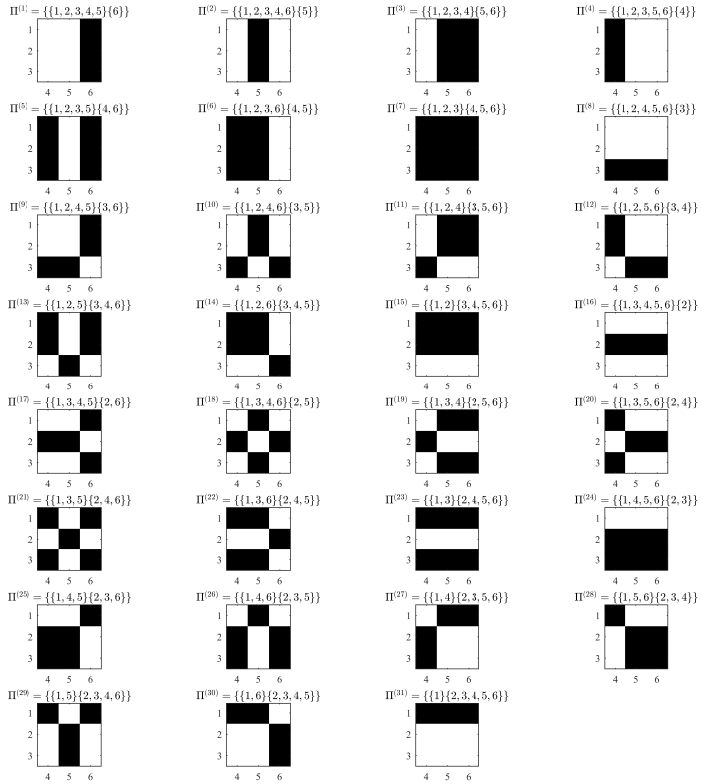

Figure S5: System decomposition-induced covariance $\mathrm{cov}(X_{t-1}, X_t)$. For a three-dimensional system, the panels depict the covariance structure induced by each possible bipartition of (X_{t-1}, X_t), here indicated by index sets (cf. eq. (50)). Black tiles indicate zero elements and white tiles indicate arbitrary elements of $\mathrm{cov}(X_{t-1}, X_t)^{(i)}$.